arXiv:1606.02900v2 [cs.PF] 8 Jan 2017

On Continuous-space Embedding of Discrete-parameter Queuing Systems

Neha V. Karanjkar Madhav P. Desai

Department of Electrical Engineering
Indian Institute of Technology Bombay
Email: nehak@ee.iitb.ac.in, madhav@ee.iitb.ac.in

Abstract

In simulation-based optimization of queuing systems, the presence of discrete-valued parameters (such as buffer sizes and the number of servers) makes the optimization difficult. We propose a novel technique for embedding such discrete parameters into a continuous space, so that optimization can be performed efficiently using continuous-space methods. Unlike spatial interpolation, our embedding technique is based on a randomization of the simulation model. The interpolated value can be computed using a single simulation of this randomized model, irrespective of the number of parameters. We first study the theoretical properties of such a randomization scheme applied to M/M/1 and Geom/Geom/1 queues. We prove that the randomization produces valid interpolations of the steady-state performance functions with respect to an integer service-time parameter. We then show that such an embedding is possible for more complex queueing networks whose parameters can include deterministic service times, queue capacities, and the number of servers, and empirically demonstrate that the interpolation is well behaved. To demonstrate the utility of the embedding technique, we solve a 6-parameter queuing network optimization problem by embedding the discrete parameters into a continuous space. The technique produces smooth interpolations of the objective function, and a continuous optimizer applied directly over this embedding converges rapidly, producing good solutions.

Keywords— *Queuing, Simulation, Discrete Optimization, Interpolation*

1 Introduction

The use of simulation is often necessary in the optimization of complex real-life queueing networks. Such queueing networks typically have discrete valued parameters such as queue capacities, server delays, the number of servers etc. More concretely, consider a queuing system with parameter set $\bar{X} = (x_1, x_2, \ldots, x_n)$, where each x_i can take an integer value. The set of all possible values that $\bar{X}$ can take is the n-dimensional discrete parameter space Ω_D. A simulation model of the system allows us to measure some objective function $f : \Omega_D \to \mathbb{R}$ (which may be composed of performance measures such as throughput, average waiting time per customer or blocking probabilities). The function f needs to be optimized. This optimization is often difficult as the number of parameters can be very large and each $f(\bar{X})$ evaluation is computationally expensive.

If the discrete parameter space can be embedded into a continuous parameter space by using some form of interpolation, the optimization problem can be solved by directly applying descent-based continuous-space techniques. Such continuous-space methods often converge fast and can scale well with the number of parameters. Such an approach has been reported in [1,2] wherein spatial interpolation was used to obtain a continuous-space embedding. However, spatial interpolation as a means of embedding is computationally expensive. For instance, in Lagrange's interpolation, evaluating the interpolated value at a single point in n dimensions would require 2^n simulations of the neighbouring discrete-space points.

We introduce a simulation-based interpolation method using which, the interpolated value can be computed in a *single* simulation run. The method is based on a randomization of the simulation model itself, and is suited to optimization problems where the objective function is composed of long-running average performance measures (such as average queue length, blocking probability, or average waiting time per customer).

The basic idea behind our interpolation scheme is to construct a randomized version of the simulation model. In this randomized model, each discrete parameter x_i in the model is *perturbed* at key points during simulation, and is assigned values of a discrete random variable X_i. The random variable X_i is chosen such that its expected value (and other moments) vary continuously with respect to some parameter $v \in \mathbb{R}$, and $X_i(v) = v$ when $v \in \mathbb{Z}$. An example of such a random

variable is:[1]

$$X_i(v) = \begin{cases} \lfloor v \rfloor & \text{with probability } \alpha(v) \\ \lceil v \rceil & \text{with probability } 1 - \alpha(v) \end{cases} \quad (1)$$

$$\text{where } \alpha(v) = \lceil v \rceil - v$$

The performance measure obtained by simulating such a randomized model is a function of v, which can now be varied continuously unlike the original discrete parameter x_i. This produces an embedding of the discrete parameter space into continuous space. As each parameter in the system can be embedded independently, the interpolated value can be obtained using a single simulation of the randomized model, irrespective of the number of parameters. We describe the randomization scheme in detail in Section 3.

To establish a basis for this randomization scheme, we study its application to an M/M/1 queue and its discrete-time counterpart, the Geom/Geom/1 queue. We first consider a Geom/Geom/1 queue with an integer service-time parameter T, where $1/T$ is the probability that the server finishes an ongoing job in the current time-slot. We introduce a randomized version of this queue, denoted as the *Interpolated Geom/Geom/1* or IGG1 queue, where the parameter T is randomized.

In the IGG1 queue, whenever the server finishes a job, the parameter T is updated and assigned values of the random variable $X_i(t)$ (defined in Equation 1). We derive expressions for the stationary distributions in the IGG1 queue and prove that its steady-state performance functions (such as average occupancy and delay) are interpolations of the corresponding functions in the original Geom/Geom/1 queue. We then extend the analysis to continuous-time versions of the two queues by treating them as a limiting case of the discrete-time queues. The analysis establishes a basis for the randomization scheme as a means of interpolation.

The basic concept of a job-based randomization of the simulation model can also be applied to parameters such as queue capacities, service time parameters and the number of servers. The embedding is not unique and the randomization can be achieved in several ways. In Section 5 we describe candidate schemes for embedding these parameters into a continuous space. We evaluate the candidate schemes using extensive simulations, and observe that they produce smooth inter-

[1] Here $\lfloor x \rfloor = \text{floor}(x)$ denotes the largest integer less than or equal to x, and $\lceil x \rceil = \text{ceiling}(x)$ denotes the smallest integer greater than or equal to x.

polations of the steady-state performance measures. The randomization scheme can be applied to both, continuous-time and discrete-time queues.

Using the randomized simulation-based embedding technique, continuous space optimization algorithms can be applied directly over the embedded model for finding optimal configurations. To demonstrate the utility of this approach, we consider a queuing network optimization problem with 6 discrete-valued parameters. We embed the parameters into a continuous space by constructing a randomized version of the model, and apply COBYLA [3], a well-established continuous-space optimization algorithm directly over the embedded model to find the solutions. We observe that the interpolation functions produced by the embedding are smooth and the optimization runs converge rapidly, producing good solutions. Compared to a conventional discrete-space search (using Adaptive Simulated Annealing) the embedding-based approach can be much more efficient. We present the results of this optimization study in Section 6.

The randomization-based embedding technique can thus be a useful tool in simulation-based optimization of complex queuing networks with a large number of discrete parameters.

2 Related Work

A survey of simulation based optimization approaches is available in [4, 5]. The discrete parameter set is typically treated separately from the continuous parameter set, using approaches such as ranking and selection, heuristic-based local search or randomized search. Discrete-parameter variants of algorithms such as SPSA [6, 7] have been used for problems where the number of dimensions is large.

If the discrete parameters have a natural integer order (for example, the number of servers), gradient-based continuous-space methods can be applied for optimization. This can be achieved in two ways. The first way is to perform optimization over a meta-model of the system [8]. A meta-model (in the form of a regression model or an artificial neural network etc.) is an approximation to the original system that is computationally cheaper to evaluate. It is constructed by systematically sampling a subset of the design space. The second way is to embed the discrete parameter space into a continuous space using some form of interpolation. Unlike a meta-model, each function evaluation of the embedded model can be as accurate

as the original simulation model. In [1] and [2], the embedding is performed using piecewise linear interpolation over a simplex. Thus interpolation in n dimensions requires $n + 1$ simulations per point.

Our work presents a simulation-based interpolation method that can produce the interpolated value in a single simulation run. It is based on a randomization of the simulation model and works by *averaging in time*, unlike spatial interpolation methods which perform an averaging in space. The technique was first used in [9] for the optimization of multi-core systems. In this paper, we give a justification for the method by analyzing its application to queuing systems, and demonstrate its utility in simulation-based optimization of queuing systems with a large number of discrete parameters.

3 Randomization Scheme

Let x_i be an integer-valued parameter in the model, and Ω_{i_D} be the set of values that x_i can take. Let $f : \Omega_{i_D} \rightarrow \mathbb{R}$ be a long-running average performance measure of the system, expressed as a function of x_i. An embedding of the parameter x_i into continuous space is essentially an interpolation of f between points in Ω_{i_D}. Let Ω_{i_C} be the set of values that the parameter can take in the embedded model. If m_i and M_i are the minimum and the maximum values in Ω_{i_D} respectively, then $\Omega_{i_C} = \{v : v \in \mathbb{R},\ m_i \leq v \leq M_i\}$. To embed the parameter x_i into continuous space, we construct an interpolation model where the value of x_i is randomized. The interpolation model can be defined at $x_i = v \in \mathbb{R}$ as follows:

Choose two integers v_1, v_2 such that v_1 is the largest integer in Ω_{i_D} less than or equal to v, and v_2 is the smallest integer in Ω_{i_D} greater than or equal to v. (If Ω_{i_D} is the set of all integers, then $v_1 = \lfloor v \rfloor$ and $v_2 = \lceil v \rceil$.) Let $h : \Omega_{i_C} \rightarrow \mathbb{R}$ be any continuous and strictly monotonic function over Ω_{i_C}. We define a function $\alpha(v)$ as:

$$\alpha(v) = \begin{cases} 0 & \text{when } v \in \Omega_{i_D} \\ \dfrac{h(v_2) - h(v)}{h(v_2) - h(v_1)} & \text{otherwise} \end{cases} \tag{2}$$

Thus $\alpha(v)$ is continuous in the interval (v_1, v_2), and $0 \leq \alpha(v) < 1$.

Let X_i be a discrete random variable with a two-point distribution

$$X_i(v,h) \;=\; \begin{cases} v_1 & \text{with probability } \alpha(v) \\ v_2 & \text{with probability } 1-\alpha(v) \end{cases} \tag{3}$$

Then in the interpolated model, the parameter x_i is assigned values of the random variable X_i at key points during simulation. Let us assume that the long-running average performance measure ($\hat{f}$) obtained by simulating this randomized model is some function of the moments of X_i. Then $\hat{f}$ is a function of v, which can now be varied continuously. This produces an embedding of the parameter x_i into continuous space. Further $\hat{f}(v) = f(v)$ when $v \in \Omega_{i_D}$. Thus $\hat{f} : \Omega_{i_C} \to \mathbb{R}$ is an interpolation of f, and $\hat{f}(v)$ can be evaluated at a given point v using a single simulation of the randomized model. The embedding is not unique and several choices exist in implementing such a randomization scheme. In particular:

- **The choice of the function** h. In most cases, we observe that choosing $h(x) = x$ or $h(x) = \frac{1}{x}$ produces a smooth interpolation.
- **The instants at which a parameter value is perturbed**. The parameter value can be perturbed once in every fixed-length time interval, or upon occurrence of a pre-defined event such as the completion of a job by a server.
- **The exact behavior of the model when a parameter value is updated**. For example, in randomizing a queue-capacity parameter, if a queue contains k entries and the queue-capacity parameter is assigned an instantaneous value of $k-1$, the excess entry in the queue can either be evicted or retained.

These choices do affect the smoothness of the generated interpolation functions. In Sections 4 and 5, we present specific examples, giving precise definitions of the randomized model's behaviour along with each example. (For all examples presented in this paper, we have used a job-based randomization. For embedding a queue-capacity parameter (Section 5.4), we assume that the excess entries in the queue are retained, and the updated parameter value is used solely for deciding whether a new entry should be allowed into the queue. Similarly, in Section 5.2 when the parameter corresponding to the number of servers at a node is updated, the jobs already in progress continue to receive service.)

4 Application of the Interpolation Scheme to Geom/Geom/1 and M/M/1 Queues

To establish a basis for the randomization-based interpolation scheme, we study its application to continuous-time and discrete-time versions of an M/M/1 queue. First, we summarize the expressions for steady-state performance measures in a discrete-time Geom/Geom/1 queue. We then introduce a job-based randomization of the Geom/Geom/1 queue (denoted the IGG1 queue). We derive expressions for the stationary distribution in the IGG1 queue and prove that the steady-state performance functions in IGG1 queue are an interpolation of the corresponding functions in the original Geom/Geom/1 with respect to a service-time parameter T. We then extend the analysis to the continuous-time versions of the two queues.

4.1 The Geom/Geom/1 Queue

A discrete-time Geom/Geom/1 queue [10,11] is defined as follows: Time is divided into equal-sized slots. Jobs arrive into the system at the beginning of a slot and depart at the end of a slot. The arrivals are Bernoulli distributed and the probability of a job arriving in a slot is p. The server has a single service-time parameter

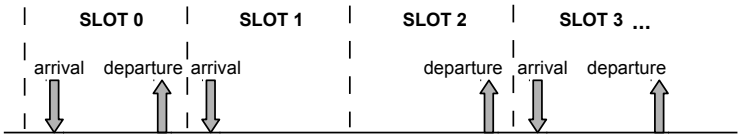

T. Whenever the queue is non-empty, the server finishes a job at the head of the queue in the current slot, with probability $q = 1/T$. The number of jobs in the system at the end of a slot can be modeled as a Markov chain in Figure 1. Let π_k be the probability that there are k jobs in the system at the end of a slot. The stationary distribution $\{\pi_k : k = 0, 1, 2, \dots\}$ can be obtained by equating the rate of transitions between states k and $k + 1$.

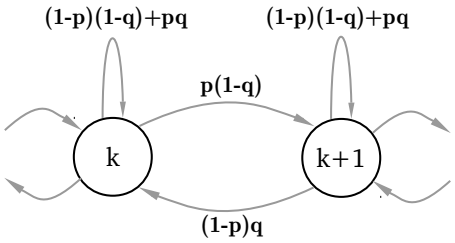

Figure 1: State transitions in a Geom/Geom/1 queue

$$\begin{aligned} p(1-q)\pi_k &= q(1-p)\pi_{k+1} \\ \pi_{k+1} &= \frac{p(1-q)}{q(1-p)}\pi_k \\ \pi_{k+1} &= \varrho\pi_k \quad \text{where } \varrho = \frac{p(1-q)}{q(1-p)} \\ \text{therefore } \pi_k &= \varrho^k\pi_0 \end{aligned}$$

Equating the sum of all state probabilities to 1, and assuming that the system is stable ($\varrho < 1$) we get:

$$\begin{aligned} 1 &= \sum_{k=0}^{\infty}\varrho^k\pi_0 \\ &= \frac{\pi_0}{(1-\varrho)} \\ \text{thus } \pi_0 &= (1-\varrho) = 1 - \frac{p(1-q)}{q(1-p)} \qquad (4) \\ \text{and } \pi_k &= \varrho^k(1-\varrho) \qquad (5) \end{aligned}$$

Let $N(T)$ be the average number of jobs in the system, and $D(T)$ be the average delay that a job experiences through the system, expressed as a function of T. Then,

$$N(T) = \sum_{k=1}^{\infty} k\pi_k \quad = \sum_{k=1}^{\infty} k\varrho^k\pi_0$$

Thus,

$$\boxed{N(T) = \frac{\varrho}{(1-\varrho)} \quad = \frac{p(1-q)}{(q-p)}, \text{ where } q = 1/T} \qquad (6)$$

Let the length of each time-slot be δ_t. The arrival rate is $\lambda = \frac{p}{\delta_t}$. Using Little's

law,

$$\boxed{D(T) = \frac{N}{\lambda} \quad = \frac{(1-q)}{(q-p)}\delta_t, \text{ where } q = 1/T} \tag{7}$$

4.2 The Interpolated Geom/Geom/1 (IGG1) Queue

The IGG1 queue is identical to a Geom/Geom/1 queue, except that the service time parameter T is itself a random variable. The behavior of the IGG1 queue can be defined with respect to a parameter $t \in \mathbb{R}_{\geq 1}$ as follows: Let $X(t)$ be a random variable with a two-point distribution:

$$X(t) = \begin{cases} \lfloor t \rfloor & \text{with probability } \alpha(t) \\ \lceil t \rceil & \text{with probability } 1 - \alpha(t) \end{cases}$$

where $\alpha(t) = \lceil t \rceil - t$. Then, each time the server finishes a job, the service-time parameter T is updated and is assigned a value X. Whenever a new job is started by the server, the probability of the server finishing the job in a cycle is the reciprocal of the service-time parameter's value.

Let $\hat{N}(t)$ denote the average system occupancy and $\hat{D}(t)$ denote the average delay of a job through the system in the IGG1 queue. $N(t)$ and $D(t)$ are the average occupancy and delay in the original Geom/Geom/1 queue given by Equations 6 and 7. We say that $\hat{N}(t)$ is an interpolation of $N(t)$ if $\hat{N}$ is continuous over $\mathbb{R}_{\geq 1}$ and $\hat{N}(t) = N(t)$ when $t \in \mathbb{Z}$. Similarly, we say that $\hat{D}(t)$ is an interpolation of $D(t)$ if $\hat{D}$ is continuous over $\mathbb{R}_{\geq 1}$ and $\hat{D}(t) = D(t)$ when $t \in \mathbb{Z}$. We prove that:

Theorem 1 *$\hat{N}$ and $\hat{D}$ are interpolations of N and D respectively.*

4.2.1 Proof:

To prove Theorem 1, we analyze the equilibrium distributions in the IGG1 queue and derive expressions for $\hat{N}$ and $\hat{D}$. The state of the system can be represented by the pair (k, T), where k is the number of jobs in the system at the end of a slot and $T \in \{\lfloor t \rfloor, \lceil t \rceil\}$ is the instantaneous value of the service time parameter. The IGG1 queue can be modeled by the Markov chain shown in Figure 2.

- For brevity, we write $\alpha(t)$ as simply α.
- Let $q_1 = \frac{1}{\lfloor t \rfloor}$, $q_2 = \frac{1}{\lceil t \rceil}$ and $q = \frac{1}{t}$. Therefore, $\frac{\alpha}{q_1} + \frac{(1-\alpha)}{q_2} = \frac{1}{q}$.
- U_k denotes the state $(k, T = \lfloor t \rfloor)$

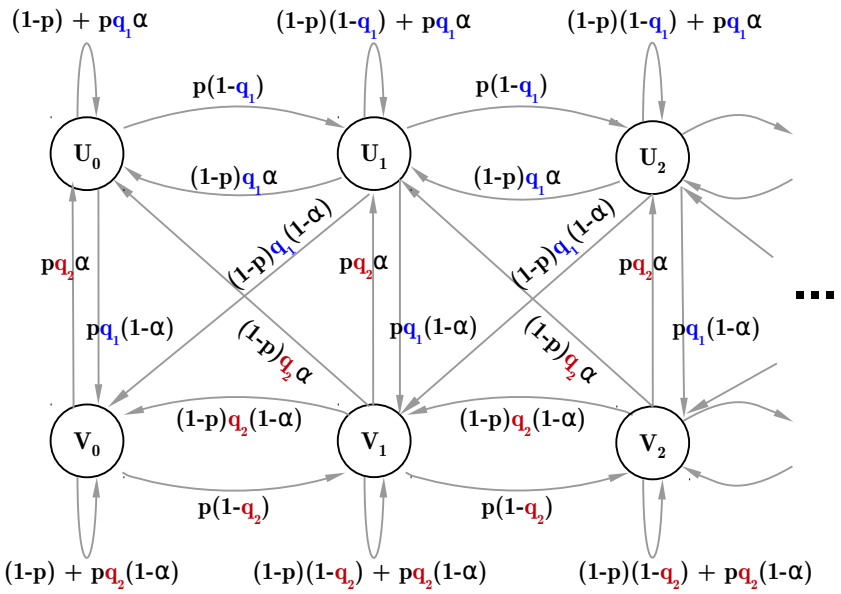

Figure 2: State transitions in the IGG1 queue

- V_k denotes the state $(k, T = \lceil t \rceil)$
- π_{U_k} and π_{V_k} are the probabilities of the system being in states U_k and V_k respectively
- $\pi_k = \pi_{U_k} + \pi_{V_k}$ is the probability that there are k jobs in the system at the end of a slot.

The Markov chain in figure 2 is irreducible and recurrent, and hence has a unique stationary distribution. It is assumed that the initial value of the service-time parameter is either $\lfloor t \rfloor$ or $\lceil t \rceil$ and this does not affect the stationary distribution. To find the distribution (π), we partition the states of the Markov chain into two groups using a *cut*. At equilibrium, the net probability of transition across such a cut should be zero. We consider four cuts as shown in Figure 3. Equating the net probability of transition across cut 1 to zero, we get
for $k \geq 0$:

$$p[(1-q_1)\pi_{U_k} + (1-q_2)\pi_{V_k}] = (1-p)[q_1\pi_{U_{k+1}} + q_2\pi_{V_{k+1}}] \tag{8}$$

By re-arranging the terms,

$$q_1\left[p\pi_{U_k} + (1-p)\pi_{U_{k+1}}\right] + q_2\left[p\pi_{V_k} + (1-p)\pi_{V_{k+1}}\right] = p(\pi_{U_k} + \pi_{V_k})$$

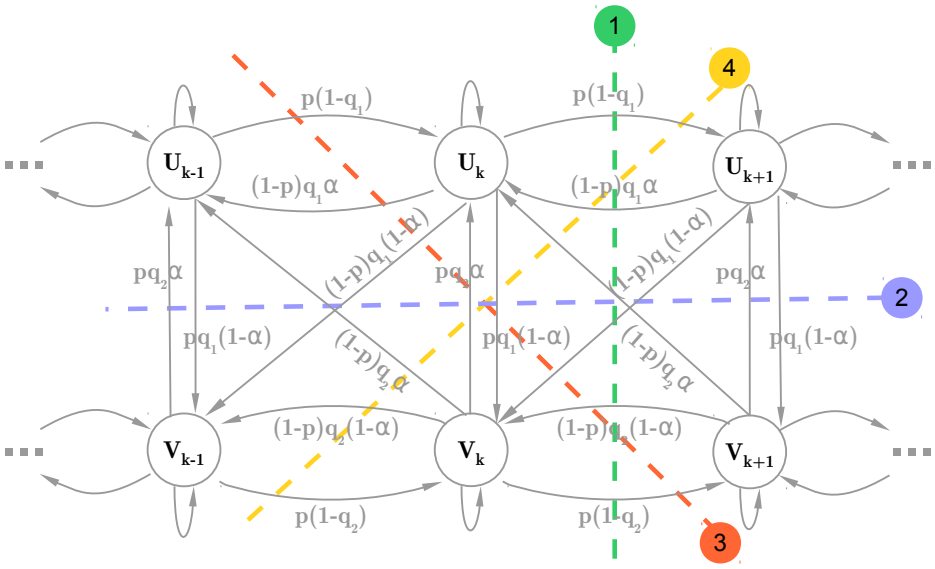

Figure 3: Cuts across the IGG1 Markov chain

Taking a summation of both sides from $k = 0$ to $k = \infty$, and substituting $\sum_{k=0}^{\infty}(\pi_{U_k} + \pi_{V_k}) = 1$, we get

$$\sum_{k=0}^{\infty} q_1 \left[p\pi_{U_k} + (1-p)\pi_{U_{k+1}}\right] + \sum_{k=0}^{\infty} q_2 \left[p\pi_{V_k} + (1-p)\pi_{V_{k+1}}\right] = p \tag{9}$$

Next, we equate the net probability of transition across cut 2 to zero

$$(1-\alpha)\sum_{k=0}^{\infty} q_1 \left[p\pi_{U_k} + (1-p)\pi_{U_{k+1}}\right] = \alpha \sum_{k=0}^{\infty} q_2 \left[p\pi_{V_k} + (1-p)\pi_{V_{k+1}}\right] \tag{10}$$

Solving Equation 9 and Equation 10 simultaneously, we get

$$\sum_{k=0}^{\infty} q_1 \left[p\pi_{U_k} + (1-p)\pi_{U_{k+1}}\right] = p\alpha \tag{11}$$

$$\sum_{k=0}^{\infty} q_2 \left[p\pi_{V_k} + (1-p)\pi_{V_{k+1}}\right] = p(1-\alpha) \tag{12}$$

Therefore

$$\sum_{k=0}^{\infty} \pi_{U_k} = \frac{p\alpha}{q_1} + (1-p)\pi_{U_0} \tag{13}$$

$$\sum_{k=0}^{\infty} \pi_{V_k} = \frac{p(1-\alpha)}{q_2} + (1-p)\pi_{V_0} \tag{14}$$

Adding Equations 13 and 14, and substituting $\sum_{k=0}^{\infty}(\pi_{U_k} + \pi_{V_k}) = 1$, we get

$$1 = p\left[\frac{\alpha}{q_1} + \frac{(1-\alpha)}{q_2}\right] + (1-p)\pi_0 \tag{15}$$

Substituting

$$\frac{\alpha}{q_1} + \frac{(1-\alpha)}{q_2} = \frac{1}{q} \tag{16}$$

we get

$$\boxed{\pi_0 = 1 - \frac{p(1-q)}{q(1-p)}} \tag{17}$$

This expression for π_0 is identical to that of the original Geom/Geom/1 queue (Equation 4). Next, we equate the net probability of transition across cuts 3 and 4 to zero.

For $k \geq 1$:

$$\begin{aligned}(pq_1\alpha - q_1)\pi_{U_k} + (p - pq_2(1-\alpha))\pi_{V_k} + p(1-q_1)\pi_{U_{k-1}} \\ = (1-\alpha)(1-p)[q_1\pi_{U_{k+1}} + q_2\pi_{V_{k+1}}]\end{aligned} \tag{18}$$

$$\begin{aligned}(p - pq_1\alpha)\pi_{U_k} + (pq_2(1-\alpha) - q_2)\pi_{V_k} + p(1-q_2)\pi_{V_{k-1}} \\ = \alpha(1-p)[q_1\pi_{U_{k+1}} + q_2\pi_{V_{k+1}}]\end{aligned} \tag{19}$$

Substituting the right-hand side of Equation 8 in Equations 18 and 19, we get

$$\left[\frac{(1-\alpha)}{(1-q_1)} + \frac{q_1(1-p)}{p(1-q_1)}\right]\pi_{U_k} + \frac{-\alpha}{(1-q_1)}\pi_{V_k} = \pi_{U_{k-1}} \tag{20}$$

$$\frac{-(1-\alpha)}{(1-q_2)}\pi_{U_k} + \left[\frac{\alpha}{(1-q_2)} + \frac{q_2(1-p)}{p(1-q_2)}\right]\pi_{V_k} = \pi_{V_{k-1}} \tag{21}$$

$$\text{Let } A = \begin{bmatrix}\left[\frac{(1-\alpha)}{(1-q_1)} + \frac{q_1(1-p)}{p(1-q_1)}\right] & \frac{-\alpha}{(1-q_1)} \\ \frac{-(1-\alpha)}{(1-q_2)} & \left[\frac{\alpha}{(1-q_2)} + \frac{q_2(1-p)}{p(1-q_2)}\right]\end{bmatrix} \tag{22}$$

$$\text{and } X_k = \begin{bmatrix}\pi_{U_k} \\ \pi_{V_k}\end{bmatrix} \tag{23}$$

Then, Equations 20 and 21 can be written in matrix form, as

$$\boxed{AX_k = X_{k-1}} \quad \text{for } k \geq 1 \tag{24}$$

To obtain X_0, we consider cuts 3 and 4 for $k = 0$. Equating the net probability of transition across these cuts to 0, we get

$$\begin{aligned}(p - pq_1\alpha)\pi_{U_0} - pq_2\alpha\pi_{V_0} &= \alpha(1-p)[q_1\pi_{U_1} + q_2\pi_{V_1}] \\ -pq_1(1-\alpha)\pi_{U_0} + (p - pq_2(1-\alpha))\pi_{V_0} &= (1-\alpha)(1-p)[q_1\pi_{U_1} + q_2\pi_{V_1}]\end{aligned}$$

Substituting the right-hand side of Equation 8 (for $k = 0$), we get

$$\pi_{U_0} = \alpha\pi_0 \tag{25}$$

$$\pi_{V_0} = (1-\alpha)\pi_0 \tag{26}$$

Thus

$$\boxed{X_0 = \begin{bmatrix}\pi_{U_0} \\ \pi_{V_0}\end{bmatrix} = \begin{bmatrix}\alpha\pi_0 \\ (1-\alpha)\pi_0\end{bmatrix}} \tag{27}$$

Equations 17, 24 and 27 together describe the stationary distribution in the IGG1 queue. The average queue occupancy ($\hat{N}$) can now be computed as

$$\hat{N} = \sum_{k=1}^{\infty}(k\pi_{U_k} + k\pi_{V_k}) \tag{28}$$

From Equation 24,

$$A\begin{bmatrix}\pi_{U_k} \\ \pi_{V_k}\end{bmatrix} = \begin{bmatrix}\pi_{U_{k-1}} \\ \pi_{V_{k-1}}\end{bmatrix} \quad \text{For } k \geq 1 \tag{29}$$

Multiplying both sides by k and taking a sum from $k = 1$ to $k = \infty$, we get

$$A\begin{bmatrix}\sum_{k=1}^{\infty} k\pi_{U_k} \\ \sum_{k=1}^{\infty} k\pi_{V_k}\end{bmatrix} = \begin{bmatrix}\sum_{k=1}^{\infty} k\pi_{U_{k-1}} \\ \sum_{k=1}^{\infty} k\pi_{V_{k-1}}\end{bmatrix}$$

This can be written as

$$A\begin{bmatrix}\sum_{k=1}^{\infty} k\pi_{U_k} \\ \sum_{k=1}^{\infty} k\pi_{V_k}\end{bmatrix} = \begin{bmatrix}\sum_{k=1}^{\infty} (k-1)\pi_{U_{k-1}} \\ \sum_{k=1}^{\infty} (k-1)\pi_{V_{k-1}}\end{bmatrix} + \begin{bmatrix}\sum_{k=1}^{\infty} \pi_{U_{k-1}} \\ \sum_{k=1}^{\infty} \pi_{V_{k-1}}\end{bmatrix}$$

Therefore

$$A\begin{bmatrix}\sum_{k=1}^{\infty} k\pi_{U_k} \\ \sum_{k=1}^{\infty} k\pi_{V_k}\end{bmatrix} = \begin{bmatrix}\sum_{k=1}^{\infty} k\pi_{U_k} \\ \sum_{k=1}^{\infty} k\pi_{V_k}\end{bmatrix} + \begin{bmatrix}\sum_{k=0}^{\infty} \pi_{U_k} \\ \sum_{k=0}^{\infty} \pi_{V_k}\end{bmatrix} \tag{30}$$

Let:

$$Y = \begin{bmatrix} \sum_{k=1}^{\infty} k\pi_{U_k} \\ \sum_{k=1}^{\infty} k\pi_{V_k} \end{bmatrix}$$
$$\text{and } S = \begin{bmatrix} \sum_{k=0}^{\infty} \pi_{U_k} \\ \sum_{k=0}^{\infty} \pi_{V_k} \end{bmatrix} \tag{31}$$

where S can be obtained from Equations 13 and 14 as:

$$S = \begin{bmatrix} \frac{p\alpha}{q_1} + (1-p)\pi_{U_0} \\ \frac{p(1-\alpha)}{q_2} + (1-p)\pi_{V_0} \end{bmatrix} \tag{32}$$

Then

$$\begin{aligned} AY &= Y + S \\ (A-I)Y &= S \\ \text{therefore } Y &= (A-I)^{-1}S \end{aligned} \tag{33}$$

where I denotes a 2×2 identity matrix. The average queue occupancy $\hat{N}$ is the sum of rows of matrix Y, and can be found by substituting A (from Equation 22) and S (from Equation 32) into equation 33. By doing these substitutions and simplifying, we obtain:

$$\hat{N} = \frac{p(1-q)}{(q-p)} + \frac{p^2 q}{(q-p)}\left[\frac{\alpha}{{q_1}^2} + \frac{(1-\alpha)}{{q_2}^2} - \frac{1}{q^2}\right] \tag{34}$$

Let δ_t be the length of each time-slot. Using Equation 16:

$$\begin{aligned} \left[\frac{\alpha}{{q_1}^2} + \frac{(1-\alpha)}{{q_2}^2} - \frac{1}{q^2}\right] &= \alpha(1-\alpha)\left[\frac{1}{q_1} - \frac{1}{q_2}\right]^2 \\ &= \alpha(1-\alpha)\left[\frac{\lfloor t \rfloor - \lceil t \rceil}{\delta_t}\right]^2 \\ &= \alpha(1-\alpha)\frac{1}{{\delta_t}^2} \end{aligned}$$

Therefore,

$$\hat{N} = \frac{p(1-q)}{(q-p)} + \frac{p^2 q}{(q-p)}\alpha(1-\alpha)\frac{1}{{\delta_t}^2} \tag{35}$$

Similarly, the average delay through the system $\hat{D}$ can be found using Little's law as $\hat{D} = \frac{\hat{N}}{p}\delta_t$. Thus,

$$\hat{D} = \frac{(1-q)}{(q-p)}\delta_t + \frac{pq}{(q-p)}\alpha(1-\alpha)\frac{1}{\delta_t} \tag{36}$$

Substituting $\delta_t = 1$ and comparing with the average occupancy N and delay D in the original Geom/Geom/1 queue from Equations 6 and 7, we get:

$$\hat{N} = N\left[1 + \frac{pq}{(1-q)}\alpha(1-\alpha)\right]$$
$$\hat{D} = D\left[1 + \frac{pq}{(1-q)}\alpha(1-\alpha)\right]$$

Expressing $\hat{N}$ and $\hat{D}$ as functions of t by substituting $q = 1/t$, we get:

$$\boxed{\begin{array}{l}\text{For } t \in \mathbb{R}_{\geq 1}: \\ \qquad \hat{N}(t) = N(t)\left[1 + \frac{p}{(t-1)}\alpha(1-\alpha)\right] \\ \qquad \hat{D}(t) = D(t)\left[1 + \frac{p}{(t-1)}\alpha(1-\alpha)\right] \\ \text{where } \alpha = (\lceil t \rceil - t).\end{array}} \tag{37}$$

Equation 37 shows that $\hat{N}$ is a continuous function of t and $\hat{N}(t) = N(t)$ when $t \in \mathbb{Z}$. Thus $\hat{N}$ is an interpolation of N. Similarly, $\hat{D}$ is an interpolation of D. This proves Theorem 1. Further, the stationary distribution π_k in the IGG1 queue is a continuous function of t, and agrees with the stationary distribution of the Geom/Geom/1 queue at integer points by definition. **Thus, the IGG1 queue represents an interpolation of the Geom/Geom/1 queue for *all* performance measures that are functions of the stationary distribution.**

4.3 The M/M/1 Queue and the Interpolated M/M/1 Queue

The M/M/1 queue can be analyzed as a limiting case of the Geom/Geom/1 queue with the length of each time-slot approaching zero. Let t be a service-time parameter of the system, λ be the arrival rate and $\mu = 1/t$ be the service-rate. The average occupancy (N_c) and the average delay of a job through the system (D_c) in the M/M/1 queue can be found by substituting $p = \lambda\delta_t$ and $q = \mu\delta_t$ in Equations 6 and 7, taking the limit $\delta_t \to 0$. Thus,

$$\varrho = \frac{\lambda}{\mu} = \lambda t \tag{38}$$
$$N_c(t) = \frac{\varrho}{(1-\varrho)} = \frac{\lambda}{(\mu-\lambda)} = \frac{\lambda t}{(1-\lambda t)} \tag{39}$$
$$D_c(t) = \frac{N}{\lambda} = \frac{1}{(\mu-\lambda)} = \frac{t}{(1-\lambda t)} \tag{40}$$

Similarly, the Interpolated M/M/1 queue can be analyzed as a limiting case of the IGG1 queue. In this queue, the service times are a mix of two hyper-exponential distributions with rates $\mu_1 = 1/\lfloor t \rfloor$ and $\mu_2 = 1/\lceil t \rceil$. The Interpolated M/M/1 queue is thus a special case of the M/H_2/1 queue [12]. The average occupancy $\hat{N}_c$ and the average delay $\hat{D}_c$ in the interpolated M/M/1 queue can be obtained from Equations 35 and 36 by substituting $p = \lambda\delta_t$ and $q = \mu\delta_t$ and taking the limit $\delta_t \to 0$:

$$\boxed{\begin{aligned} &\text{For } t \in \mathbb{R}_{>1} : \\ &\qquad \hat{N}_c(t) = N_c(t)\left[1 + \frac{\lambda}{t}\ \alpha(1-\alpha)\right] \\ &\qquad \hat{D}_c(t) = D_c(t)\left[1 + \frac{\lambda}{t}\ \alpha(1-\alpha)\right] \\ &\quad \text{where } \alpha = (\lceil t \rceil - t). \end{aligned}} \tag{41}$$

Thus $\hat{N}_c(t)$ and $\hat{D}_c(t)$ are interpolations of $N_c(t)$ and $D_c(t)$ respectively. The relative error between the performance measures in the original and the interpolated queues is

$$\frac{\hat{N}_c(t) - N_c(t)}{N_c(t)} = \frac{\hat{D}_c(t) - D_c(t)}{D_c(t)} = \frac{\lambda}{t}\,\alpha(1-\alpha) \tag{42}$$

In Figure 4 we plot this error term as a function of t. The error term diminishes as the size of the perturbation relative to the parameter value ($\frac{(\lceil t \rceil - \lfloor t \rfloor)}{t}$) becomes smaller and smaller, which is intuitive. The error is bounded below 0.125 (the maximum error occurs at $t = 1.\overline{3}$). Thus, not only are $\hat{N}$ and $\hat{D}$ valid interpolations, but closely approximate N and D.

We have shown that the randomization scheme produces valid interpolations of the steady-state performance functions (such as average occupancy and delay) in both the discrete-time and continuous-time versions of the M/M/1 queue with respect to the service-time parameter. Further, the randomized queue closely approximates the original queue when the perturbations relative to the parameter value are small.

5 Interpolation in More Complex Queuing Systems

We show that the randomization-based embedding strategy can be extended to parameters such as queue capacities, deterministic service times and number of

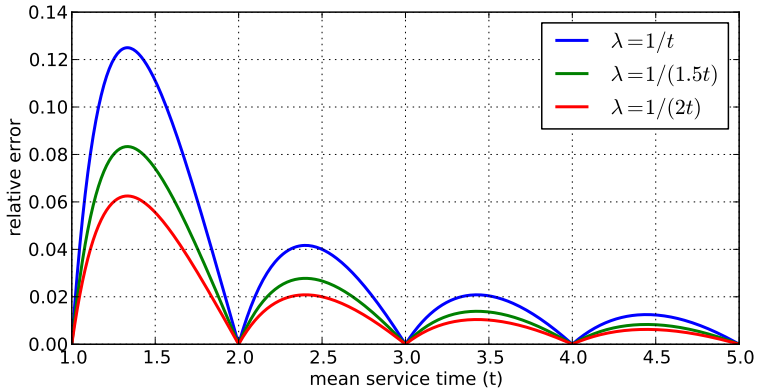

Figure 4: Relative error (as defined by Equation 42) between the original and the Interpolated M/M/1 queues

servers. We consider several examples of queuing systems and evaluate the interpolation using simulations of the randomized (embedded) model alongside the original (discrete-parameter) model. Although we report simulation results for discrete-time queues, the same randomization scheme can be applied to corresponding continuous-time queues (as illustrated by the analysis in the previous section).

5.1 Embedding service time (T) in a Geom/D/1 queue

Consider a Geom/D/1 queue (a discrete time counterpart of an M/D/1 queue). Arrivals are Bernoulli distributed and the probability of a job arriving in a slot is p. The server is deterministic. Each job accepted by the server takes a constant time of T slots to finish, where $T \in \mathbb{N}$. To embed the discrete parameter T into continuous space, we construct an interpolated model where T is randomized. To evaluate the interpolation at $v \in \mathbb{R}$, we choose a function $h(x) = x$ and obtain the random variable $X(v, h)$ (as introduced in Eq. 3). In the simulation, whenever a job is accepted by the server, we assign the job a service time of X slots.

In Figure 5, we show the average number of jobs in the system ($N(T)$) observed in a simulation of the interpolated Geom/D/1 queue (with $p = 0.1$). The service time parameter T is swept in steps of 0.1 and $N(T)$ is computed at each point using a single simulation run of the interpolated model. We see that the in-

terpolated function $N(T)$ is continuous and smooth and agrees with the original Geom/D/1 queue at integer values of T. Thus the randomization produces an observably smooth embedding of the discrete parameter T.

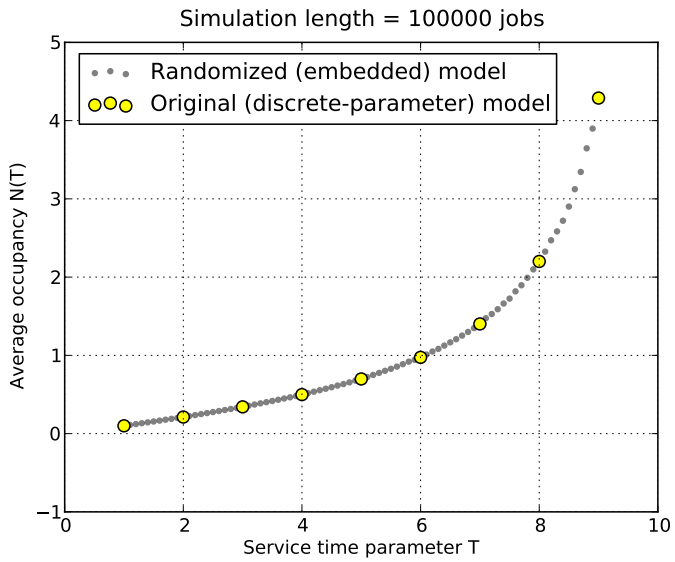

Figure 5: Mean number of jobs in the system (N) versus service time (T) in a Geom/D/1 queue ($p = 0.1$)

5.2 Embedding the number of servers (K) in Geom/Geom/K and Geom/D/K queues

Consider a Geom/Geom/K queue (a discrete-time counterpart of an M/M/K queue). The system consists of an infinite-capacity queue and K identical servers working in parallel. The probability of a job arriving into the queue in a time-slot is p. The service times are geometrically distributed, and the probability of each server finishing a job in a time-slot is q.

We wish to embed the integer parameter K into continuous space. To do so, we first assert that a system with K identical, independent servers can be described (or modelled) in the following alternative manner:

An alternative description:

The system consists of a single infinitely fast *controller* with a parameter K, and a fixed large number ($> K$) of identical, independent servers. The behavior of the controller with respect to the parameter K is as follows:

- In each time-slot, the controller pulls jobs from the head of the queue and assigns each job to a free server, as long as the queue is not empty and the number of jobs currently receiving service is less than K.
- The controller then iterates over the list of ongoing jobs. For each ongoing job, the job is either terminated in the current slot with probability q, or continues to receive service with probability $(1 - q)$.

This alternative description allows us to present the randomization scheme in a more intuitive manner, now in terms of the controller parameter K. To embed the discrete parameter K into continuous space, we construct an interpolated model by randomizing K. To evaluate the interpolation at $v \in \mathbb{R}$, we choose $h(x) = x$ and a random variable $X(v, h)$ (as introduced in Eq. 3). Then, in the interpolated model, whenever the controller pulls a new job from the queue, the parameter K is modified by setting it to the random variable X. By this definition, whenever the parameter K is updated, the jobs already receiving service are not disturbed and the updated parameter value is used solely for deciding if a new job can be pulled from the queue to start receiving service.

In Figure 6, we show the simulation results obtained with this interpolation scheme for a Geom/Geom/K queue with arrival probability $p = 1/10$ and service probability $q = 1/19$. The interpolated function $N(K)$ is the average number of jobs in the system as a function of K.

The interpolation scheme for a Geom/D/K queue is identical to that for a Geom/Geom/K queue, except that all jobs started by the server take T slots to finish, where T is a constant. In Figure 7, we show the simulation results obtained with this interpolation scheme for a Geom/D/K queue with arrival probability $p = 1/10$ and service time $T = 19$ slots. The interpolated function $N(K)$ is the average number of jobs in the system as a function of K. The plots indicate that the randomization scheme produces smooth embeddings of the discrete parameter K.

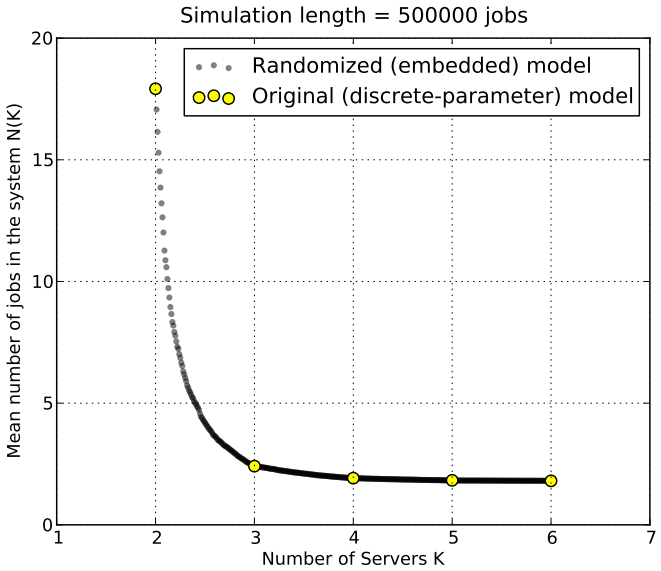

Figure 6: Mean number of jobs in the system (N) versus number of servers (K) in a Geom/Geom/K queue ($p = \frac{1}{10}, q = \frac{1}{19}$)

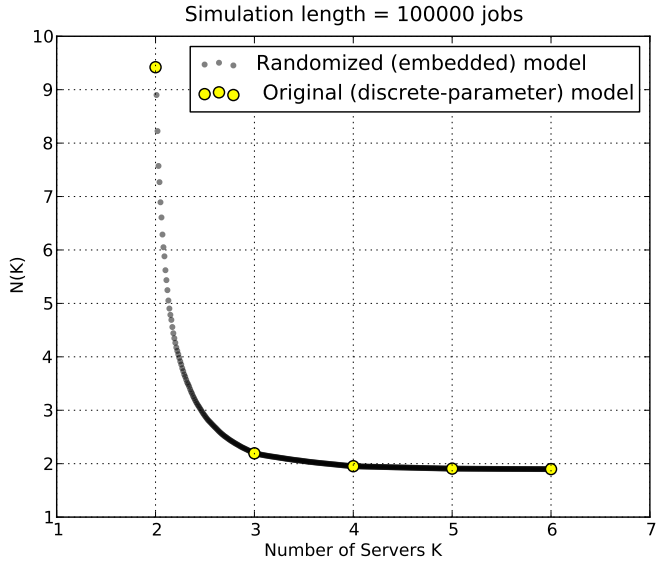

Figure 7: Mean number of jobs in the system (N) versus number of servers (K) in a Geom/D/K queue ($p = \frac{1}{10}, T = 19$)

5.3 Interpolation-based embedding of multiple parameters: the number of servers (K) and the service time (T) in a Geom/D/K queue

For a discrete-time Geom/D/K queue, with K identical servers working in parallel and a deterministic service time T, the behavior of the system can be defined in terms of a single controller as follows:

- In each time-slot, the controller pulls jobs from the head of the queue as long as the queue is not empty and the number of jobs currently receiving service is less than K. Each job j pulled by the controller is assigned a service time $T_j = T$.
- The controller then iterates over the list of ongoing jobs, and removes jobs that have expired. A job j expires after T_j slots since its start of service.

In the interpolated model the parameters T and K are randomized. To evaluate the interpolation at some point $(v_T, v_K) \in \mathbb{R}^2$ in the continuous parameter space, we choose $h(x) = x$, and two independent random variables $X_1(v_T, h)$ and $X_2(v_K, h)$ (as defined in Equation 3). Whenever the server accepts a new job from the queue, the parameter T is modified by setting it to the random variable X_1, and the parameter K is set to the random variable X_2.

In Figure 8, we show the simulation results obtained with this interpolation scheme for a Geom/D/K queue with arrival probability $p = 1/10$. The interpolated function $N(T, K)$ is the average number of jobs in the system. The parameter T is swept in steps of 0.1 and K is swept in steps of 0.05. At each point (T, K) a single simulation run of the interpolated model is used to obtain $N(T, K)$. We observe that the interpolation surface is continuous and smooth. This example illustrates that multiple parameters in the model can be embedded simultaneously.

5.4 Embedding queue capacity (C) in a Geom/Geom/1/C+1 queue

Consider a Geom/Geom/1/C+1 queue (the discrete-time counterpart of an M/M/1/C+1 queue). The system is identical to the Geom/Geom/1 queue, except that the queue in front of the server has a finite capacity. Let C denote the queue-capacity parameter. We define the behaviour of the queue with respect to this parameter C (where $C \in \mathbb{N}$) as follows:

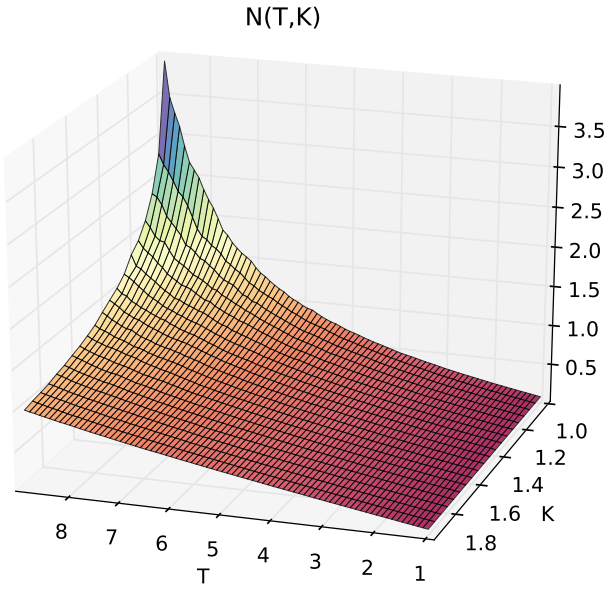

Figure 8: Mean number of jobs in the system (N) as a function of the number of servers (K) and the service time (T) in a Geom/D/K queue ($p = \frac{1}{10}$, Simulation length = 500000 jobs)

- Whenever a job arrives, if the number of jobs already present in the queue is less than the value of parameter C, the job is allowed to enter the queue. Else, the job is lost.

To embed the parameter C into continuous space, we construct an interpolated Geom/Geom/1/C+1 queue by randomizing C. To evaluate the interpolation at $v \in \mathbb{R}$, we choose $h(x) = x$ and a random variable $X(v, h)$. Whenever a job arrives and tries to enter the queue, the parameter C is modified by setting it to the random variable X. By this definition, whenever the parameter value is updated, the jobs already present in the queue are not disturbed. The updated parameter value is used solely for deciding if the next arriving job should be allowed to enter the queue. (Thus it is possible that the occupancy of the queue at some point of time can be greater than the instantaneous value of C.)

In Figure 9, we show the simulation results obtained with this randomization scheme for a system with arrival probability $p = 1/10$ and service probability $q = 1/9$. The queue capacity is swept between $(1, 9)$ in steps of 0.1. The interpolated

function $P_B(C)$ is the blocking probability (probability that an arriving job is lost). We observe that the randomization scheme produces a smooth interpolation of the steady-state blocking probability function.

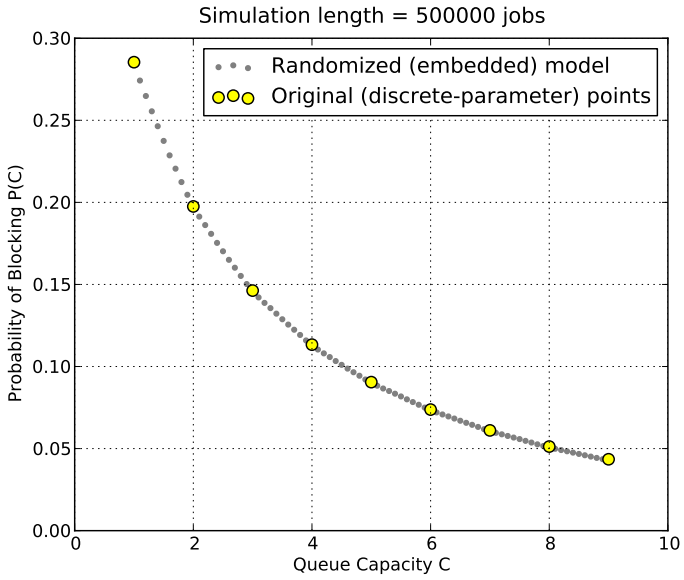

Figure 9: Probability of Blocking (P_B) versus queue capacity (C) in a Geom/Geom/1/C+1 queue ($p = \frac{1}{10}, q = \frac{1}{9}$)

5.5 Embedding queue capacity (C) in a Geom/D/1/C+1 queue

The interpolation scheme for a Geom/D/1/C+1 queue is identical to that for a Geom/Geom/1/C+1 queue (described in Section 5.4), except that all jobs started by the server take T slots to finish, where T is a constant. The quantity being interpolated is the blocking probability (P_B) as a function of queue capacity C.

We observe that choosing $h(x) = x$ produces an interpolation with kinks (Figure 10a). A smoother interpolation can be obtained by choosing $h(x) = \frac{1}{x}$ (Figure 10b). Thus the choice of the randomization scheme can affect the smoothness of the interpolated function, and needs to be made with care.

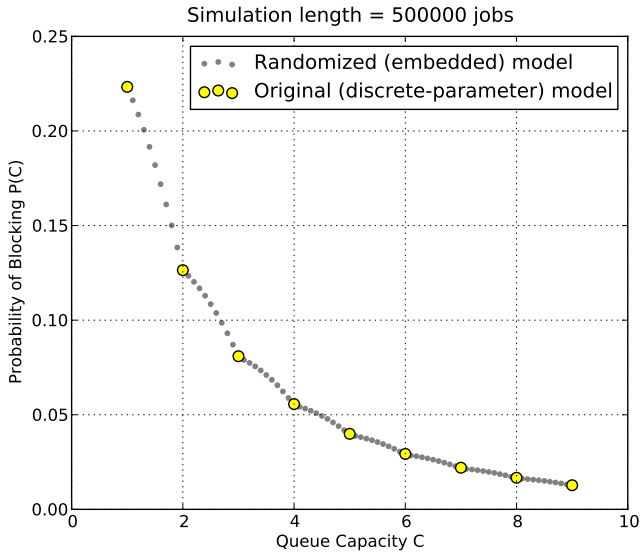

(a) Interpolation produced using $h(x) = x$

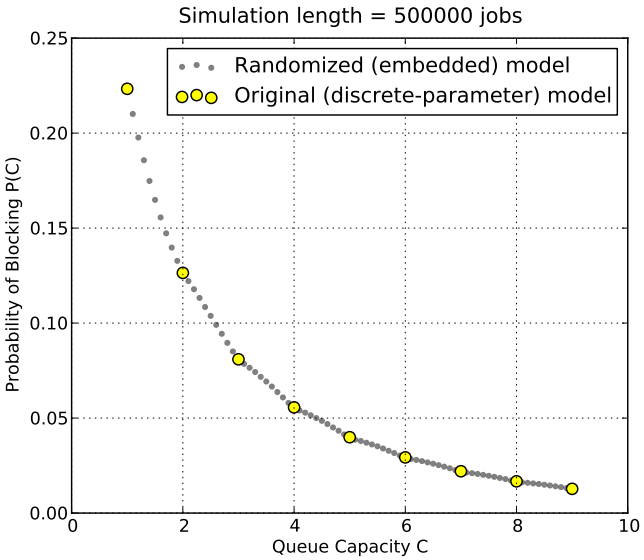

(b) Interpolation produced using $h(x) = \frac{1}{x}$

Figure 10: Probability of Blocking (P_B) versus queue capacity (C) in a Geom/D/1/C+1 queue ($p = \frac{1}{10}, T = 9$)

6 An Optimization Example

Using the randomized simulation-based interpolation scheme, the discrete parameters in a queuing system can be embedded into a continuous space. Optimization can then be performed using continuous-space techniques, applied directly over the interpolated model. To demonstrate this approach we consider the optimization of a queuing network shown in Figure 11:

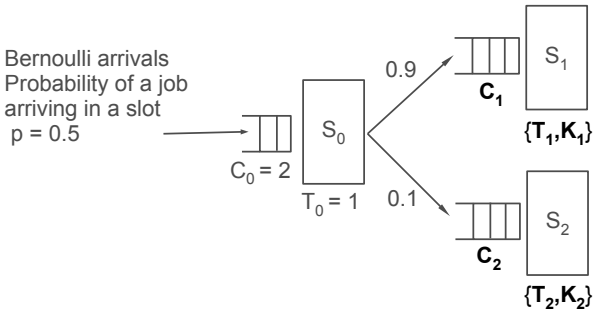

Figure 11: Queuing system being optimized

- The system consists of three servers S_0, S_1 and S_2 with deterministic service times and a finite capacity queue in front of each server. Each server i has service time T_i and a queue capacity C_i.
- The servers S_1 and S_2 internally consist of multiple identical servers working in parallel. The number of servers in S_1 and S_2 are K_1 and K_2 respectively.
- Arrivals into the system are Bernoulli distributed. The probability of a job arriving in a slot is p. Upon arrival of a job, if the queue in front of server S_0 is full, the job is lost.
- Upon processing each job, the server S_0 forwards the job to either S_1 (with probability 0.9) or to S_2 (with probability 0.1). If the destination queue selected for a job is full, the server S_0 blocks and does not process any jobs until the current job can be placed into the destination queue.

Optimization is performed over the parameter set $\{C_1, T_1, K_1, C_2, T_2, K_2\}$. The other parameters are kept fixed at $p = 0.5, C_0 = 2$ and $T_0 = 1$. The function

to be minimized (f) is a weighted sum of cost and throughput components. Cost is defined as :

$$\text{cost} = \frac{K_1}{T_1} + \frac{K_2}{T_2}$$

The objective function is :

$$f = W \times \text{cost} - \text{throughput}$$

where W is a constant representing the weight assigned to cost during optimization. Table 1 lists the bounds and constraints on the parameter values. The parameters are embedded into a continuous space by constructing a randomized interpolation model. The randomization schemes used for embedding service times (T_1, T_2) and the number of servers (K_1, K_2) are identical to those described in Section 5. The queues in front of servers S_1 and S_2 are fed by a blocking source. For embedding the queue capacity parameters (C_1, C_2), we have used a randomization with $h(x) = x$.

Table 1: Parameter bounds and constraints for optimization

Parameter	Range (min,max)	Parameter	Range(min,max)
C_1	(1 ,4)	C_2	(1 ,4)
T_1	(10 ,∞)	T_2	(10 ,∞)
K_1	(1 ,10)	K_2	(1 ,10)
Constraints : $K_1 + K_2 \leq 10,\ C_1 + C_2 \leq 4$			

Over the interpolated model, we apply COBYLA [3], a continuous-space optimization algorithm that does not require computation of gradients, and can accept linear constraints. We use an existing implementation of COBYLA from Python's SciPy library [13]. For each evaluation of f, the simulation length was chosen to be 50,000 jobs. On an average, each optimization run makes about 60 function evaluations and takes 2-3 seconds to execute on a desktop PC. We perform several optimization runs by sweeping the weight W assigned to cost, and using multiple randomly chosen initial points, to get a cost-performance trade-off curve (Figure 12). Each point in the plot represents the results of a single optimization run, with the (x, y) coordinates of the point representing (cost, throughput) values at the optimum. The plot shows a clear *knee*. The solutions are well-clustered and

all solutions near the knee region are observed to have the ratio K_1/T_1 close to 0.5 and K_2/T_2 close to 0.1.

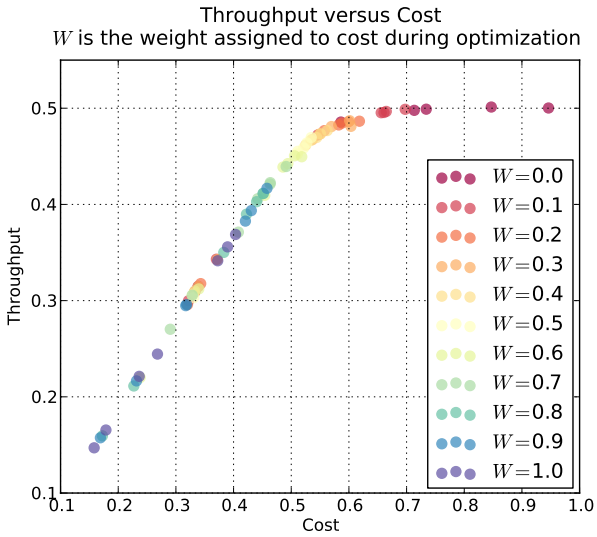

Figure 12: Pareto front obtained by sweeping W. Each point represents the results of a single optimization run, with the (x, y) coordinates representing (cost, throughput) at the optimum

To assess the smoothness of the interpolation, we plot the surface of the objective function f near an optimum ($(C_1, T_1, K_1, C_2, T_2, K_2) = (3, 17, 8, 1, 20, 2)$) with $W = 0.2$. As the function is 6-dimensional, we plot f with respect to two parameters at a time by keeping the other parameter values fixed at their optimum (Figure 13). The plots indicate that near the optimum, the interpolation is smooth.

To highlight the advantage of this embedding technique, we compare the performance of the embedding-based optimization approach against a discrete-space search performed directly over the original (discrete-parameter) model. For the discrete-space search, we use a well-established implementation of the Adaptive Simulated Annealing (ASA) algorithm [14, 15]. The optimization problem for this comparative study is same as the 6-parameter problem defined in Section 6 with a weight-factor $W = 0.2$, and without the constraints $K_1 + K_2 \leq 10,\ C_1 + C_2 \leq 4$.[2] We first solve this problem by directly performing a discrete-space search using ASA. A single simulation of the original discrete-parameter model is used for each

[2]The ASA implementation accepts only box constraints.

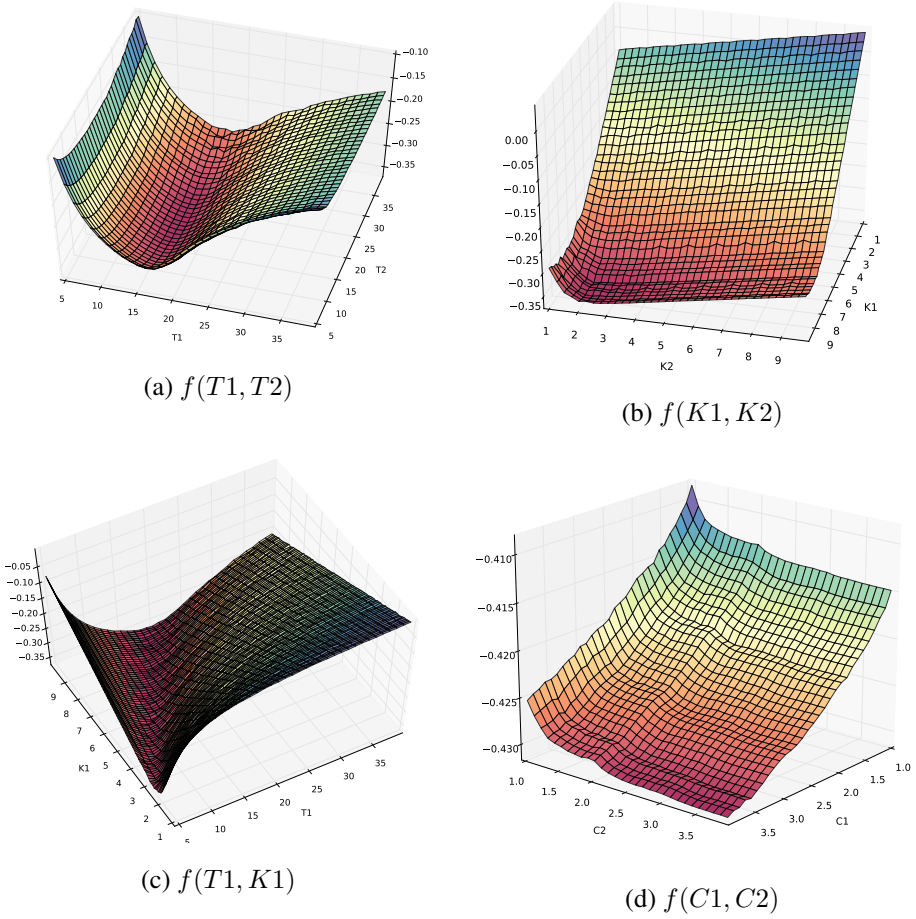

(a) $f(T1, T2)$

(b) $f(K1, K2)$

(c) $f(T1, K1)$

(d) $f(C1, C2)$

Figure 13: Surface plots of the objective function f near the optimum $(C_1, T_1, K_1, C_2, T_2, K_2) = (3, 17, 8, 1, 20, 2)$ illustrating that the interpolation functions produced using the randomization-scheme are smooth.

objective evaluation during optimization. The length of each simulation run is $1,00,000$ jobs. We perform ten such optimization runs with distinct (random) initial guesses and find the average number of objective evaluations per optimization run. We similarly perform ten optimization runs of COBYLA applied over the embedded (randomized) model and compare the performance of the two approaches side-by-side in Table 2. The results indicate that an optimization performed using the embedding-based approach can be much more efficient (faster) than a direct discrete-space search, while producing solutions of comparable quality. This optimization example illustrated the utility of the continuous-space embedding tech-

nique.

Table 2: Performance comparison between a discrete-space search (ASA) over the original model and a continuous-space optimization (COBYLA) over the embedded (randomized) model

		ASA (original model)	COBYLA (embedded model)
Number of simulations per optimization run	mean	1964	69
	std-dev	344	10
Time per optimization run in seconds	mean	137.0	4.9
	std-dev	25.3	0.9
Objective-value at optimum	best	-0.383	-0.383
	mean	-0.382	-0.370
	std-dev	0.0001	0.0160

7 Conclusions

We have presented a simple, efficient technique for embedding discrete parameters in a queuing system into continuous space so that continuous space methods can be used to optimize the queuing system. The technique is based on a randomization of the simulation model, and can produce the interpolated value in a single simulation run. We observe that the technique produces good embeddings for parameters such as queue capacities and the number of servers using several examples.

To demonstrate the utility of this technique, we have solved a queuing network optimization problem with 6 parameters. The technique produced a smooth embedding of the discrete parameter space, and a continuous optimization algorithm can be effectively used on the embedding. The randomized simulation-based interpolation method can thus be a useful tool in the optimization of complex discrete parameter queuing networks.

References

[1] H. Wang and B. W. Schmeiser, "Discrete Stochastic Optimization using Linear Interpolation," in *2008 Winter Simulation Conference*. IEEE, Dec. 2008.

[2] E. Lim, "Stochastic Approximation over Multidimensional Discrete Sets with Applications to Inventory Systems and Admission Control of Queueing Networks," *ACM Trans. Model. Comput. Simul.*, vol. 22, no. 4, pp. 19:1–19:23, Nov. 2012.

[3] M. Powell, "On Trust Region Methods for Unconstrained Minimization without Derivatives," *Mathematical Programming*, vol. 97, no. 3, 2003.

[4] M. C. Fu, "Optimization via simulation: A review," *Annals of Operations Research*, vol. 53, no. 1, pp. 199–247, 1994.

[5] J. R. Swisher, P. D. Hyden, S. H. Jacobson, and L. W. Schruben, "A Survey of Simulation Optimization Techniques and Procedures," in *Simulation Conference, 2000. Proceedings. Winter*, vol. 1, 2000, pp. 119–128 vol.1.

[6] J. E. Whitney, L. I. Solomon, and S. D. Hill, "Constrained optimization over discrete sets via SPSA with application to non-separable resource allocation," in *Proceeding of the 2001 Winter Simulation Conference (Cat. No.01CH37304)*, vol. 1, 2001, pp. 313–317 vol.1.

[7] V. Mishra, S. Bhatnagar, and N. Hemachandra, "Discrete parameter simulation optimization algorithms with applications to admission control with dependent service times," in *2007 46th IEEE Conference on Decision and Control*, Dec 2007, pp. 2986–2991.

[8] R. R. Barton and M. Meckesheimer, "Chapter 18: Metamodel-Based Simulation Optimization ," in *Simulation*, ser. Handbooks in Operations Research and Management Science, S. G. Henderson and B. L. Nelson, Eds. Elsevier, 2006, vol. 13, pp. 535 – 574.

[9] N. V. Karanjkar and M. P. Desai, "An Approach to Discrete Parameter Design Space Exploration of Multi-core Systems Using a Novel Simulation Based Interpolation Technique," in *Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), 2015 IEEE 23rd International Symposium on*, Oct 2015, pp. 85–88.

[10] T. G. Robertazzi, *Networks and Grids: Technology and Theory*, 1st ed. Springer Publishing Company, Incorporated, 2010.

[11] H. Bruneel, "Performance of discrete-time queueing systems," *Computers & Operations Research*, vol. 20, no. 3, pp. 303–320, 1993.

[12] D. R. Insua, M. Wiper, and F. Ruggeri, "Bayesian analysis of M/Er/1 and M/H_k/1 queues," *Queueing Systems*, vol. 30, no. 3, pp. 289–308, 1998.

[13] Implementation of COBYLA in Python's scipy.optimize library. [Online]. Available: http://docs.scipy.org/doc/scipy-0.14.0/reference/generated/scipy.optimize.fmin_cobyla.html

[14] J. Robert. Python bindings for the ASA code. [Online]. Available: https://pypi.python.org/pypi/pyasa/

[15] L. Ingber, "Adaptive Simulated Annealing (ASA): Lessons Learned," *Control and Cybernetics*, vol. 25, 1996.